

Digital Data System:

User's View of the Network

By J. J. MAHONEY, JR., J. J. MANSELL, and R. C. MATLACK

(Manuscript received July 12, 1974)

The utilitarian aspect of the Digital Data System is discussed with emphasis on performance objectives that will be important when data communications is inserted into a system of data processing. Objectives for the dependability and quality of data communications are quantified and evaluated in terms of their impact on data processing. Characteristics of the several types of channels available are described in detail along with operational features of particular importance at the interface between data communication and data processing.

I. PERFORMANCE OBJECTIVES

1.1 General

Modern industries are motivated to centralize the control of their integrated operations at a powerful computer and to disperse important operational functions to outlying areas. Modern telecommunications provide the data transmission services that are essential to coordinate dispersed functions with centralized control. Viewed in this light, data communication becomes one link in a larger chain of operations. Objectives for the reliability and quality of data transmission are evaluated in terms of their impact on the larger operation. Data service should generally be deemed excellent when the capability of a data processing system is not diminished perceptibly by the insertion of data communication.

1.2 Availability

The dependability of data service may be expressed in terms of availability, which is simply the complement of average annual down time. The dependability of the service of the DDS is quite naturally related to the inherent reliability of the network of telecommunications facilities from which individual data transmission paths are derived. Dependability of service has been enhanced by providing for auto-

matic substitution of standby facilities during emergencies. The most vulnerable part of the service tends to be the portion of plant that is necessarily dedicated to a single user, where redundant facilities would be most expensive. With automatic substitution to maintain continuity of service during failure of facilities in the common network and with prompt repair of dedicated plant, the design objective for DNS services is to attain average availability of at least 99.96 percent. This objective would permit average annual down time no greater than 210 minutes, which amounts to 0.04 percent of the total time in the 24-hour days of a 365-day year. Thus, communications down time is expected to be a small fraction of the total down time of an entire system of data processing.

The causes of communication failure cover such a wide range of possibilities that it is not possible to specify an amount of down time expected in a particular year for the maximum duration of failure if it should occur.

1.3 Quality

Along with stringent objectives for availability and maintainability, a quantitative description of the quality of data communications is necessary to guide efficient use of the service. It has been traditional to describe the quality of data communication in a simple one-number characterization called error rate, which is the ratio of bit errors to the total number of bits transmitted in a test period. Efficient utilization of data services, however, depends on the exact structure of error patterns as well as the average probability of bit error. For example, where data is blocked into messages containing substantial numbers of bits, the throughput efficiency realized is sensitive to the correlation among bit errors. If bit errors are independent (i.e., no correlation), the number of messages that must be retransmitted to avoid error is very nearly equal to the number of bit errors. Where bit errors are clumped within the block interval, the number of error-free messages is substantially greater than indicated by the average probability of bit error. To minimize the impact of imperfect data communication on the overall system of data processing, the design objective for average efficiency of data communication is to attain 99.5-percent error-free seconds. The percentage of error-free seconds is not related to the more familiar bit-error rate in any simple way, but the proportion of error-free seconds is often estimated by the mathematically convenient relationship,

$$P(\text{EFS}) = (1 - P_e)^B,$$

which gives the probability of an error-free second that would arise

from the probability of bit-error P_e at bit rate B . This formula consistently underestimates the percentage of error-free seconds because it relies on complete independence among bit-errors, an assumption known to be invalid. In a practical communications network, transmission is subject to occasional and momentary perturbations, e.g., the automatic substitution of standby facilities to avoid outage of significant duration. Such perturbations seldom persist longer than a small fraction of a second, but they may be interpreted as error bursts. Short-term unavailability of this type is included in allowances of erroneous seconds. Thus, the objective is properly stated in terms of error-free seconds without implication about the ratio of bits found in error to total bits transmitted.

The 1-second sample is short enough to display high resolution in the description of data communication quality, yet long enough to encompass most block lengths that users may choose for error control, and long enough to span most momentary degradations of service. Thus, 1-second samples should approach the behavior of independent events in probabilistic analysis. Since 1 second is regarded as an upper bound on the duration of error bursts, the user may attain transmission efficiency greater than the objective for data communications efficiency if his block length is less than the number of bits transmitted in 1 second. The throughput he actually attains will, of course, depend on irreducible transit time in communication over long distances and on other delays, including those encountered in his own method of error control.

1.4 Allocation of quality

A model of the limiting nns connection would comprise a long-haul network interconnecting two local serving areas in which there are two local T1 carrier lines in tandem with one baseband loop. The results of field tests, modified by experienced judgment, lead to allocating the total tolerance of errors in the following way:

- 2 baseband loops would account for 2/10
- 4 local lines would account for 3/10
- 1 long-haul network connection would account for 5/10.

The carrier facilities (both long-haul and local) are interfaced at the DS-1 bit rate of 1.544 Mb/s, and the error performance objectives are expressed in error-free seconds at the DS-1 interface for those facilities. The method of allocation is based on the general formula:

$$\text{Percent EFS objective} = 100 - TAP/N,$$

where

T = total tolerance of 0.5 percent, the complement of 99.5-percent EFS

A = fractional allocation

P = number of ports demultiplexed

N = an empirical factor relating the average number of errored seconds expected in DS-0 ports for each one found on the facility.

Appropriate N factors are not yet firmly established but estimates derived from preliminary field tests indicate that 99.6-percent EFS would be applicable to individual local carrier lines. On a comparable basis, the indicated objective would be 99-percent EFS for long-haul network connections. Both objectives would apply at the DS-1 interface where the bit rate is 1.544 Mb/s.

1.5 Performance estimates

Maintainability studies are continuing, but they are necessarily based on hypothetical probabilities of failure and estimates of the time that will be required to restore service. Contributions to annual down time from many of the possible failures have been rendered negligible by providing for automatic substitution of standby facilities. Where automatic protection of service is not provided, system fault-location features have been built in to reduce the duration of service outage by minimizing the time required to locate faults. Study results to date indicate that the objective for availability will be feasible for the large majority of connections expected. Preliminary results of field tests of individual subsystems indicate that the objective for quality of data communications will generally be met.

II. POINT-TO-POINT DDS SERVICE

The duplex private line types of channels are point-to-point and multipoint. As indicated in the article by Snow and Knapp, these are synchronous channels operating at 2.4, 4.8, 9.6, or 56 kb/s. Two choices of terminations are available to the customer. The first is known as the data service unit (DSU) and the second is the channel service unit (CSU) which is part of the basic channel offering.

2.1 Data Service Unit (DSU)

The DSU is physically located on the customer's premises with the output connected to a four-wire loop facility connected to the central

office and the input connected to the customer's data-terminal equipment. The DSU consists of two basic sections, a channel terminator, and an encoder-decoder. The function of the channel terminator is (i) to provide a balanced termination for the four-wire loops, and (ii) to provide the circuitry for implementing the loopback tests. The encoder-decoder consists of a transmitter, a receiver, and a clock recovery section and provides the EIA and CCITT drivers and terminators that interface with the data terminal equipment. The basic function of this unit is the conversion of EIA RS-232-C or CCITT V.35 interface signals to baseband bipolar line signals, and vice versa.

2.2 Data terminal interface

The DSU uses one of two interface connectors, depending on the service offering, one for 2.4-, 4.8-, or 9.6-kb/s service, and the other for 56-kb/s service. For the former, the interface signals exchanged between the data terminal and the DSU are in bipolar voltage form and conform to EIA Standard RS-232-C.

For 56-kb/s service, the clocks and data (transmit and receive) are dc coupled balanced signals. These signals and the interface circuits involved meet the balanced interface standard of CCITT Recommendation V.35. The control signals conform to EIA Standard RS-232-C. For 56-kb/s service, all interface signals should be transmitted over balanced-pair conductors for improved performance and less crosstalk.

For each service offering the interface circuits provided by the DSU match the Type-D interface of RS-232-C for dedicated line service. When the Permanent Request to Send option is used, the DSU has a Type-E interface. Therefore, the DSU provides "plug-for-plug" interchangeability with Type-D or Type-E interfaces of present data sets used on private-line, analog, voiceband networks.

2.3 System operation

The DMS, in the initial service offering, provides for two-point, four-wire, duplex, private-line, digital-data transmission. Although the DMS is a four-wire network, the customer terminals may operate either in a one-way, half-duplex, or in a duplex manner. In describing the operations of the DSU, four modes of operation can be defined: data, idle, out of service, and test. The transmitting section of the DSU can attain the data or idle mode independent of the state of the receiving section, while the receiving section can attain the data, idle, or out of service mode independent of the state of the transmitter section. The test mode involves both the transmitting and receiving section of the DSU.

2.3.1 Half-duplex operation

In half-duplex operation only one terminal transmits at a time. The data terminal desiring to transmit switches its request-to-send circuit on. After a delay, the clear-to-send circuit switches on, indicating that the data terminal may begin transmission. The receiving data terminal has its request-to-send circuit switched off.

To turn the circuit around, the transmitting data-terminal equipment should send an end-of-message (EOM) code and then switch its request-to-send circuit off. Upon receiving the EOM code, the receiving data terminal equipment switches the request-to-send circuit on and after a short delay receives a clear-to-send on signal. If the permanent-request-to-send option is used, the receiving terminal may start transmitting immediately after the EOM code is received. The transmission delay between terminals consists of the propagation delay determined by the routing of the specific circuit and a fixed delay through Bell System terminal equipments. The transmission delay for one-way transmissions over terrestrial facilities will generally be less than 50 ms.

For transmit-only service, it is advisable that the permanent-request-to-send option be used to avoid the clear-to-send delay.

2.3.2 Duplex operation

Since the DDS provides four-wire point-to-point service, and simultaneous transmission in both directions is possible, it is convenient to use the permanent request-to-send option of the DSU so that the clear-to-send circuit is always on. With this option, the data terminal equipment must have a Type-E interface of EIA RS-232-C. When the request-to-send circuit is under the control of the data terminal equipment, the DSU has a Type-D interface.

2.4 Testing and maintenance

The DSU provides testing ability under manual-switch control on the DSU or under the control of the Serving Test Center (STC). When the DSU is in the Test mode, an indication is given to the customer by means of either the LL (local test) lamp or the RT (remote test) lamp.

2.4.1 Manual control of test modes

A test switch provides the customer with the capability of performing an LL or an RT.

2.4.1.1 Local test. With the test switch in the LL position, the DSU is in the local test mode. The LL test permits the customer with a duplex terminal to test the back-to-back performance of his data-terminal equipment and DSU by connecting the transmitter section of the DSU

to the receiver section. In addition, the receive line is connected through terminating equipment to the transmit line to allow a signal to be maintained in both directions. For this test the data-set-ready circuit is switched off, but the other control interface circuits, request-to-send, clear-to-send, and received-line-signal-detector, operate as in the idle or data mode.

When the LL test switch is operated, the line is looped in both directions. This gives the remote terminal the capability of testing the transmission path to and from the local DSU as well as permitting the local terminal to test its DSU, as described above. During LL test, the clock of the local DSU is held in synchronization with the system clock.

2.4.1.2 Remote test. With the test switch in the RT position, the DSU is in the remote test mode. In this test mode the output of the received data circuit is connected to the input of the transmitted data circuit at the data terminal interface of the DSU. For this test, the control interface circuit drivers to the data-terminal equipment are switched off and the transmitted data and received data circuits from and to the customer are left open.

With the local DSU in the RT test mode, the remote data terminal has the capability of checking system operation exclusive of the local data terminal. This permits the customer to deduce whether the local data terminal is responsible for a system trouble condition.

2.4.2 Remote control of test modes from the serving test center

In addition to the manual control of the test modes, the telephone company's STC can place the DSU in either the LL or RT mode to test the operations of the line and DSU.

In the RT mode, the STC ascertains whether there are any defects in the transmitter, receiver, and interface circuits of the DSU and the transmission path to and from the customer. It does not ascertain whether the customer is putting proper signals on the interface circuits.

If the results of the RT test show that there is a trouble condition, then the STC can place the DSU in the LL test mode to isolate the trouble condition between the DSU and the transmission path.

2.5 Channel service unit (CSU)

An optional interface to the DSU is available and is known as the CSU. It provides a minimum channel termination that allows for the remote testing of the local DDS channel.

Nominal 50-percent duty cycle, bipolar pulses are accepted from the customer on the data transmit (DT) and data receive (DR) leads. These pulses must be synchronous with the DDS and limited to a specified maximum jitter. The input bipolar pulses are amplified and

filtered, and pass through the transmit-repeat coil to the transmit pair. The received signal is first amplified and equalized and then sliced. The resultant bipolar pulses are then passed to the customer. From these pulses, the customer must recover the synchronous clock used for timing the transmit data and sampling the received data.

III. MULTIPOINT DDS SERVICE

Multipoint service has existed in data communications from the earliest days of telegraphy as means for sharing a single channel among several stations which individually do not generate sufficient traffic to justify a full-time dedicated circuit. With multipoint DDS service, three or more customer stations may be connected onto a single circuit that can be operated more efficiently than existing circuits due to lower channel turnaround delays. DDS makes multipoint service available at 2.4 kb/s, 4.8 kb/s, and, for the first time, 9.6 kb/s and 56 kb/s. The customer stations may be located at a number of different customer sites served by DDS and all stations on a single circuit will transmit at the same DDS transmission rate (i.e., 2.4 kb/s, 4.8 kb/s, 9.6 kb/s, or 56 kb/s). All DDS multipoint circuits have two-way simultaneous transmission capability, but can, of course, be used in a half-duplex or one-way-at-a-time fashion. Like two-point DDS circuits, multipoint DDS circuits are transparent (no coding restrictions) to the content and format of the data which the customer transmits.

DDS multipoint service serves only those multipoint circuits that have a single customer-control location and a number of outlying stations. Communication is from any outlying station to the control location and from the control location to any set of outlying stations. Communication from one outlying station to another is not possible except via the control location. All data transmitted by the control location is delivered by DDS to every station on the multipoint circuit.

The customer is responsible for the overall supervision and signaling control of the circuit. His responsibilities include:

- (i) Inserting addressing information at the control station to permit outlying stations to determine if the information is destined for them.
- (ii) Detecting the addressing information at outlying stations.
- (iii) Supervising the communications circuit from the control location to insure that outlying stations do not attempt to transmit simultaneously.

These responsibilities are normally performed by the customer through the use of a communications controller at outlying stations and a computer at the control location.

3.1 Multipoint junction unit

Multipoint service capability in the DDS is provided by interconnecting standard, point-to-point customer channels at the 64-kb/s level by means of Multipoint Junction Units (MJU). To facilitate testing, MJUs are located in DDS hub offices. Although multipoint service is available at all DDS standard data rates, the MJU itself is independent of customer data rates since input and output is at the 64-kb/s level.

MJUs are inserted into a multipoint circuit at hub locations, where they can split the data path from the control station into two or more branches directed toward the outlying stations and combine the data branches from the outlying station into one path toward the control station. The MJU itself is a two-circuit card device that can provide one path toward the control location and two branches on one card, or one path toward the control location and up to four branches on a total of two cards. Additional branches are obtained by cascading MJUs with a cascade of N four-branch MJUs resulting in $3N + 1$ available branches.

Data being transmitted along the communications path from the control location is passed through the MJU unaltered and delivered to all branches toward the outlying stations. Data transmitted from outlying stations to the control location enter a 1-byte serial-shift register whose output is connected to the input of an AND gate. If an incoming byte from a branch is a data byte, it passes unchanged to the AND-gate input. If the byte is a network-control byte, however, it is changed to the all-1's byte when passed to the AND gate. In effect, the control byte is suppressed to a data byte of all 1's. This results in confining possible trouble condition indications to a single branch, thus preventing interference with communication from other branches.

3.2 Circuit operation and implications

There are at least two major ways of operating a DDS multistation line. In the first method, all stations, including the control station, normally have the request-to-send (RTS) circuit to the DSU switched off. The control station begins operation by switching the RTS on and waiting for the clear-to-send (CTS) circuit from the DSU to be switched on. When the CTS is switched on, the control station begins transmitting addressing characters which are to be delivered to all outlying stations via the digital channels and MJUs. These addressing characters indicate to customer-supplied communication controllers at outlying stations whether the data message is destined for them. The station selected then switches on its RTS lead to its DSU, waits for CTS from its DSU, and responds to selection with a message to the control station. Communi-

cation then continues between these two as if they were on a two-point private channel. When communication is completed, the control station and outlying station exchange completion notices and switch off their RTS leads, and the circuit returns to its idle state.

The implications of this type of operation are that the idle state of the multipoint nns circuit has circuit-idle characters generated by nns office channel units (OCUs) in all channels. When the control station is transmitting or receiving from an outlying station, data is flowing along those channels but idle codes remain in all other branches. The idle codes, however, are suppressed by the MJUs in the circuit as described previously, thus there is no interference from other branches. Moreover if transmission errors occur on idle branches, the errors are suppressed unless they happen to convert the idle code into a data byte. The advantage of this type of operation has a corresponding disadvantage, namely, that this security has been paid for by sacrificing the time required to start up communications (i.e., switching RTS to on, waiting for CTS, etc.).

An alternative method for operating a multipoint nns circuit is to require all stations (including the control station) to normally have their RTS lead switched on and to be transmitting a constant steady series of 1's as an idle pattern. In this case, there is no need to wait for CTS, and the control location can begin operation by merely transmitting addressing characters. As described above, all stations receive these characters, one is selected, and that one may immediately begin transmitting to the control station. Of course, the data transmitted from this station is combined by the MJU with data from all other branches; however, the other stations are transmitting a steady series of 1's and, when they are combined with data in the ANN gate of the MJU, the data remains as it was on input.

This type of operation clearly has less delay and, consequently, higher efficiency than the one described previously since there is no requirement to switch RTS on and wait for CTS before transmitting. However, the inherent disadvantage is that to the MJU, it appears that all stations are transmitting data at all times, even though most stations are transmitting an idle pattern of steady 1's. Therefore, any transmission errors in any branch of the multipoint circuit, which convert a 1 to a 0, will cause errors to occur and propagate through the circuit to the control location.

The nns multipoint is designed to be compatible with either type of operation described above. The customer chooses by which method he wishes to operate on the basis of the amount of delay he can tolerate and the degree of protection against transmission errors he desires.

3.3 Delay and efficiency of transmission

An important attribute of any multipoint service is the absolute time delay introduced by the communications system. This delay is important because the basic motivation behind multipoint service is one of providing an economical means for connecting several lightly loaded stations to a single circuit. Obviously, wasting more time in establishing a connection on a multipoint circuit leaves less time available for data communication, and, consequently, fewer stations may be connected to a single circuit.




Fig. 1—Typical polling response delays for DDS circuits.

The basic elements of the DDS that introduce delay are the rate at which electrical signals propagate through the transmission media and the built in bit and byte delays in DSUs, OCU's, MJUs, and digital multiplexing equipment. For any specific circuit, these delays can be accurately predicted from knowledge of the distances involved, the equipment through which the circuit is routed, and the actual transmission speed of the circuit.

Figure 1 provides an indication of polling response delay for DDS multipoint circuits as a function of the distance between the control station and the outlying station. This figure is based on the assumption that the customer's master or control station continually keeps its RTS circuit switched on, and transmits polling characters to the remote station. The remote station then switches its RTS on, waits for a CTS indication from the DSU, then transmits a response. Therefore, the delays encountered are one round-trip propagation delay plus one start-up delay at the remote DSU.

When using existing analog data sets, the time delay between RTS and CTS in itself generally exceeds the delays indicated on Fig. 1 by significant amounts. It should be clear, therefore, that similar curves for analog circuits would generally show DDS to have less delay and, consequently, DDS would permit the user to operate his circuit more efficiently than he could using analog. However, the user must be aware that these curves are representative of typical delays and do not portray what may be encountered in any individual case. The specific delays encountered are dependent on the exact routing of a particular channel. Moreover, the user should be aware of the possibility that the delays he experiences on an individual circuit may change due to internal network rearrangements or due to network trouble conditions.

IV. ACKNOWLEDGMENTS

Original contributors to the information in Section II, which describes operation of the Data Service Unit and the Channel Service Unit, are W. D. Farmer, J. R. Klepper, and A. L. Pappas.